

All-but-the-Top: Simple and Effective Postprocessing for Word Representations

Jiaqi Mu, Suma Bhat, and Pramod Viswanath

University of Illinois at Urbana Champaign

{jiaqimu2, spbhat2, pramodv}@illinois.edu

Abstract

Real-valued word representations have transformed NLP applications; popular examples are word2vec and GloVe, recognized for their ability to capture linguistic regularities. In this paper, we demonstrate a *very simple*, and yet counter-intuitive, postprocessing technique – eliminate the common mean vector and a few top dominating directions from the word vectors – that renders off-the-shelf representations *even stronger*. The postprocessing is empirically validated on a variety of lexical-level intrinsic tasks (word similarity, concept categorization, word analogy) and sentence-level extrinsic tasks (semantic textual similarity) on multiple datasets and with a variety of representation methods and hyperparameter choices in multiple languages; in each case, the processed representations are consistently better than the original ones. Furthermore, we demonstrate quantitatively in downstream applications that neural network architectures “automatically learn” the post-processing operation.

1 Introduction

Words and their interactions (as sentences) are the basic units of natural language. Although words are readily modeled as discrete atomic units, this is unable to capture the relation between the words. Recent distributional real-valued representations of words (examples: word2vec, GloVe) have transformed the landscape of NLP applications – for instance, text classification (Socher et al., 2013b; Maas et al., 2011; Kim, 2014), machine translation (Sutskever et al., 2014; Bahdanau et al., 2014) and knowledge base completion (Bor-

des et al., 2013; Socher et al., 2013a). The success comes from the geometry of the representations that efficiently captures linguistic regularities: the semantic similarity of words are well captured by the similarity of the corresponding vector representations; the latent relation between two words is well captured by the difference vector between the corresponding two vector representations.

A variety of approaches have been proposed in recent years to learn the word representations: Collobert et al. (2011); Turian et al. (2010) learn the representations via semi-supervised learning by jointly training the language model and downstream applications; Bengio et al. (2003); Mikolov et al. (2010); Huang et al. (2012) do so by fitting the data into a neural network language model; Mikolov et al. (2013); Mnih and Hinton (2007) by log-linear models; and Dhillon et al. (2012); Pennington et al. (2014); Levy and Goldberg (2014); Stratos et al. (2015); Arora et al. (2016) by producing a low-dimensional representation of the cooccurrence statistics. Despite the wide disparity of algorithms to induce word representations, the performance of several of the recent methods is roughly similar on a variety of intrinsic and extrinsic evaluation testbeds.

In this paper, we find that a *simple* processing renders the off-the-shelf existing representations *even stronger*. The proposed algorithm is motivated by the following observation.

Observation *Every representation we tested, in many languages, has the following properties:*

- The word representations have *non-zero mean* – indeed, word vectors share a large common vector (with norm up to a half of the average norm of word vector).
- After removing the common mean vector, the representations are *far from* isotropic – in-

deed, much of the energy of most word vectors is contained in a very low dimensional subspace (say, 8 dimensions out of 300).

Implication Since all words share the same common vector and have the same dominating directions, and such vector and directions strongly influence the word representations in the same way, we propose to eliminate them by: (a) removing the nonzero mean vector from all word vectors, effectively reducing the energy; (b) projecting the representations *away* from the dominating D directions, effectively reducing the dimension. Experiments suggest that D depends on the representations (for example, the dimension of the representation, the training methods and their specific hyperparameters, the training corpus) and also depends on the downstream applications. Nevertheless, a rule of thumb of choosing D around $d/100$, where d is the dimension of the word representations, works uniformly well across multiple languages and multiple representations and multiple test scenarios.

We emphasize that the proposed postprocessing is *counter intuitive* – typically denoising by dimensionality reduction is done by eliminating the *weakest* directions (in a singular value decomposition of the stacked word vectors), and *not* the dominating ones. Yet, such postprocessing yields a “purified” word representation as seen in our elaborate experiments.

Experiments By postprocessing the word representation by eliminating the common parts, we find the processed word representations to capture stronger linguistic regularities. We demonstrate this quantitatively, by comparing the performance of both the original word representations and the processed ones on three canonical lexical-level tasks:

- *word similarity* task tests the extent to which the representations capture the similarity between two words – the processed representations are consistently better on seven different datasets, on average by 2.3%;
- *concept categorization* task tests the extent to which the clusters of word representations capture the word semantics – the processed representations are consistently better on three different datasets, by 2.8%, 4.5% and 4.3%;

- *word analogy* task tests the extent to which the difference of two representations captures a latent linguistic relation – again, the performance is consistently improved (by 0.5% on semantic analogies, 0.2% on syntactic analogies and 0.4% in total). Since part of the dominant components are inherently canceled due to the subtraction operation while solving the analogy, we posit that the performance improvement is not as pronounced as earlier.

Extrinsic evaluations provide a way to test the goodness of representations in specific downstream tasks. We evaluate the effect of postprocessing on a standardized and important extrinsic evaluation task on sentence modeling: *semantic textual similarity* task – where we represent a sentence by its averaged word vectors and score the similarity between a pair of sentences by the cosine similarity between the corresponding sentence representation. postprocessing improves the performance consistently and significantly over 21 different datasets (average improvement of 4%).

Word representations have been particularly successful in NLP applications involving supervised-learning, especially in conjunction with neural network architectures. Nodes in all neural network architectures have enough flexibility to conduct the postprocessing step (i.e., subtract a bias term and null out a subspace) and thus one conjectures that perhaps postprocessing does not significantly affect the end-to-end performance. Indeed, we see this in an experiment on a standard *text classification* task on five datasets using a well established convolutional neural network (CNN) classifier (Kim, 2014). The performance improves only mildly with postprocessing over all five datasets.

Towards verifying our hypothesis that the neural network “effectively learns to postprocess” within its nodes, we conduct the following experiment. We append an extra *linear* layer to the beginning of a recurrent neural network (RNN) where the classifier is built on the last hidden unit of the RNN. The bias term of the linear layer is jointly trained with the parameters of the RNN to best fit the training set (of the same five different datasets of the text classification task). We find that the bias term of the linear layer when inferred with the postprocessed word representations plus the mean vector of word representations is *very close*

to the bias term of the linear layer when inferred with the original word representations (the cosine similarity is roughly 0.69, across the five datasets and three different RNN architectures – the vanilla RNN, GRU-RNN (Chung et al., 2015) and LSTM-RNN (Greff et al., 2016)). This allows us to conclude that the neural networks “learn” to postprocess the word representations.

We summarize our contributions below:

- We observe a novel geometric property of word representations on every representation learning algorithm we tested;
- We propose a *very simple* postprocessing algorithm to purify the word representations so that they can capture stronger linguistic regularities;
- We quantitatively demonstrate the superiority of the processed word representations using intrinsic and extrinsic evaluations, over a very wide range of standardized datasets, word representation methods with varying hyperparameters and in multiple languages.

Related Work Our work is directly related to word representation algorithms, most of which have been elaborately cited in the introduction and also throughout the entire paper, at the appropriate locations. To the best of our knowledge, this is the first paper discussing a postprocessing algorithm on *pre-trained* word representations.

2 Postprocessing

We test our observations on various word representations: four publicly available word representations (WORD2VEC¹ (Mikolov et al., 2013) trained using Google News, GLOVE² (Pennington et al., 2014) trained using Common Crawl, RAND-WALK (Arora et al., 2016) trained using Wikipedia and TSCCA³ trained using English Gigaword) and two self-trained word representations using CBOW and Skip-gram (Mikolov et al., 2013) on the 2010 Wikipedia corpus from (Al-Rfou et al., 2013). For completeness, we also consider the representations on other languages: we use the publicly available TSCCA representations (Dhillon et al., 2012) on German, French, Spanish,

Italian, Dutch and Chinese. The detailed statistics for all representations are listed in Table 1.

Let $v(w) \in \mathbb{R}^d$ be a word representation for a given word w in the vocabulary V . We observe the following two phenomena in each of the word representations listed above:

- $\{v(w) : w \in V\}$ are not of zero-mean: i.e., all $v(w)$ share a non-zero common vector,

$$v(w) = \tilde{v}(w) + \mu,$$

where μ is the average of all $v(w)$ ’s, i.e., $\mu = 1/|V| \sum_{w \in V} v(w)$. As reported in Table 1, the norm of μ is approximately 1/6 to 1/2 of the average norm of all $v(w)$.

- $\{\tilde{v}(w) : w \in V\}$ are not isotropic: Let u_1, \dots, u_d be the first to the last components recovered by the principal component analysis (PCA) of $\{\tilde{v}(w) : w \in V\}$, and $\sigma_1, \dots, \sigma_d$ be the corresponding normalized variance ratio. Each $\tilde{v}(w)$ can be written as a linear combinations of u :

$$\tilde{v}(w) = \sum_{i=1}^d \alpha_i(w) u_i.$$

As shown in Figure 1, we observe that σ_i decays near exponentially for small values of i and remains roughly constant over the later ones. This suggests there exists D such that $\alpha_i \gg \alpha_j$ for all $i \leq D$ and $j \gg D$; from Figure 1 one observes that D is roughly 10 with dimension $d = 300$.

Since all word representations share the same common vector μ and have the same dominating directions and such vector and directions strongly influence the word representations in the same way, we propose to eliminate them, as formally achieved in Algorithm 1.

Discussion In our proposed processing algorithm, the number of components to be nulled, D , is the only hyperparameter that needs to be tuned. We find that a good rule of thumb is to choose D approximately to be $d/100$, where d is the dimension of a word representation. This is empirically justified in the experiments of the following section where $d = 300$ is standard for published word representations. We trained word representations for higher values of d using the WORD2VEC and GLOVE algorithms and repeat these experiments; we see corresponding consistent improvements due to postprocessing in Appendix A.


¹<https://code.google.com/archive/p/word2vec/>

²<https://github.com/stanfordnlp/GloVe>


³<http://www.pdhillon.com/code.html>

	Language	Corpus	dim	vocab size	avg. $\ v(w)\ _2$	$\ \mu\ _2$
WORD2VEC	English	Google News	300	3,000,000	2.04	0.69
GLOVE	English	Common Crawl	300	2,196,017	8.30	3.15
RAND-WALK	English	Wikipedia	300	68,430	2.27	0.70
CBOW	English	Wikipedia	300	1,028,961	1.14	0.29
Skip-Gram	English	Wikipedia	300	1,028,961	2.32	1.25
TSCCA-En	English	Gigawords	200	300,000	4.38	0.78
TSCCA-De	German	Newswire	200	300,000	4.52	0.79
TSCCA-Fr	French	Gigaword	200	300,000	4.34	0.81
TSCCA-Es	Spanish	Gigaword	200	300,000	4.17	0.79
TSCCA-It	Italian	Newswire+Wiki	200	300,000	4.34	0.79
TSCCA-Nl	Dutch	Newswire+Wiki	200	300,000	4.46	0.72
TSCCA-Zh	Chinese	Gigaword	200	300,000	4.51	0.89

Table 1: A detailed description for the embeddings in this paper.



(a) English



(b) Multilingual

Figure 1: The decay of the normalized singular values of word representation.

3 Experiments

Given the popularity and widespread use of WORD2VEC (Mikolov et al., 2013) and GLOVE (Pennington et al., 2014), we use their publicly available pre-trained representations in the following experiments. We choose $D = 3$ for WORD2VEC and $D = 2$ for GLOVE.

Algorithm 1: All-But-The-Top: a postprocessing algorithm on word representations.

Input : Word representations

$\{v(w), w \in V\}$, a threshold
parameter D ,

1 Compute the mean of $\{v(w), w \in V\}$,

$$\mu \leftarrow \frac{1}{|V|} \sum_{w \in V} v(w); \quad \tilde{v}(w) \leftarrow v(w) - \mu$$

2 Compute the PCA components:

$$u_1, \dots, u_d \leftarrow \text{PCA}(\{\tilde{v}(w), w \in V\}).$$

3 postprocess the representations:

$$v'(w) \leftarrow \tilde{v} - \sum_{i=1}^D (u_i^T \tilde{v}(w)) u_i$$

Output: Processed representations $v'(w)$.

The key underlying principle behind word representations is that similar words should have similar representations. Following the tradition of evaluating word representations (Schnabel et al., 2015; Baroni et al., 2014), we perform three canonical *lexical-level* tasks: (a) word similarity; (b) concept categorization; (c) word analogy; and one sentence-level task: (d) semantic textual similarity. The processed representations consistently improve performance on all three of them, and especially strongly on the first two.

3.1 Word Similarity

The word similarity task is as follows: given a pair of words, the algorithm assigns a “similarity” score – if the pair of words are highly related then the score should also be high and vice versa. The algorithm is evaluated in terms of Spearman’s rank correlation compared to (a gold set of) human judgements.

For this experiment, we use seven standard datasets: the first published RG65 dataset (Rubenstein and Goodenough, 1965); the widely used WordSim-353 (WS) dataset (Finkelstein et al., 2001) which contains 353 pairs of commonly used verbs and nouns; the rare-words (RW) dataset (Luo et al., 2013) composed of rarely used words; the MEN dataset (Bruni et al., 2014) where the 3000 pairs of words are rated by crowdsourced participants; the MTurk dataset (Radinsky et al., 2011) where the 287 pairs of words are rated in terms of relatedness; the SimLex-999 (SimLex) dataset (Hill et al., 2016) where the score measures “genuine” similarity; and lastly the SimVerb-3500 (SimVerb) dataset (Gerz et al., 2016), a newly released large dataset focusing on similarity of verbs.

In our experiment, the algorithm scores the similarity between two words by the cosine similarity between the two corresponding word vectors ($\text{CosSim}(v_1, v_2) = v_1^T v_2 / \|v_1\| \|v_2\|$). The detailed performance on the seven datasets is reported in Table 2, where we see a consistent and significant performance improvement due to post-processing, across all seven datasets. These statistics (average improvement of 2.3%) suggest that by removing the common parts, the remaining word representations are able to capture stronger semantic relatedness/similarity between words.

	WORD2VEC		GLOVE	
	orig.	proc.	orig.	proc.
RG65	76.08	78.34	76.96	74.36
WS	68.29	69.05	73.79	76.79
RW	53.74	54.33	46.41	52.04
MEN	78.20	79.08	80.49	81.78
MTurk	68.23	69.35	69.29	70.85
SimLex	44.20	45.10	40.83	44.97
SimVerb	36.35	36.50	28.33	32.23

Table 2: Before-After results (x100) on word similarity task on seven datasets.

3.2 Concept Categorization

This task is an indirect evaluation of the similarity principle: given a set of concepts, the algorithm needs to group them into different categories (for example, “bear” and “cat” are both animals and “city” and “country” are both related to districts). The clustering performance is then evaluated in terms of purity (Manning et al., 2008) – the fraction of the total number of the objects that were classified correctly.

We conduct this task on three different datasets: the Almuhareb-Poesio (ap) dataset (Almuhareb, 2006) contains 402 concepts which fall into 21 categories; the ESSLLI 2008 Distributional Semantic Workshop shared-task dataset (Baroni et al., 2008) that contains 44 concepts in 6 categories; and the Battig test set (Baroni and Lenci, 2010) that contains 83 words in 10 categories.

Here we follow the setting and the proposed algorithm in (Baroni et al., 2014; Schnabel et al., 2015) – we cluster words (via their representations) using the classical k -Means algorithm (with fixed k). Again, the processed vectors perform consistently better on all three datasets (with average improvement of 2.5%); the full details are in Table 3.

	WORD2VEC		GLOVE	
	orig.	proc.	orig.	proc.
ap	54.43	57.72	64.18	65.42
esslli	75.00	84.09	81.82	81.82
battig	71.97	81.71	86.59	86.59

Table 3: Before-After results (x100) on the categorization task.

3.3 Word Analogy

The analogy task tests to what extent the word representations can encode latent linguistic relations between a pair of words. Given three words w_1 , w_2 , and w_3 , the analogy task requires the algorithm to find the word w_4 such that w_4 is to w_3 as w_2 is to w_1 .

We use the analogy dataset introduced in (Mikolov et al., 2013). The dataset can be divided into two parts: (a) the *semantic* part containing around 9k questions, focusing on the latent semantic relation between pairs of words (for example, what is to Chicago as Texas is to Houston); and (b) the *syntactic* one containing roughly 10.5k questions, focusing on the latent syntactic relation

between pairs of words (for example, what is to “amazing” as “apprently” is to “apparent”).

In our setting, we use the original algorithm introduced in (Mikolov et al., 2013) to solve this problem, i.e., w_4 is the word that maximize the cosine similarity between $v(w_4)$ and $v(w_2) - v(w_1) + v(w_3)$.

The detailed performance on the analogy task is provided in Table 4. It can be noticed that while postprocessing continues to improve the performance, the improvement is not as pronounced as earlier. We hypothesize that this is because the mean and some dominant components get canceled during the subtraction of $v(w_2)$ from $v(w_1)$, and therefore the effect of postprocessing is less relevant.

	WORD2VEC		GLOVE	
	orig.	proc.	orig.	proc.
syn.	73.46	73.50	74.95	75.40
sem.	72.28	73.36	79.22	79.25
all	72.93	73.44	76.89	77.15

Table 4: Before-After results (x100) on the word analogy task.

3.4 Semantic Textual Similarity

Extrinsic evaluations measure the contribution of a word representation to specific downstream tasks; below, we study the effect of postprocessing on a standard sentence modeling task – *semantic textual similarity* which aims at testing the degree to which the algorithm can capture the semantic equivalence between two sentences. For each pair of sentences, the algorithm needs to measure how similar the two sentences are. The degree to which this measure matches with human judgment (in terms of Pearson’s correlation) is an index of the algorithm’s performance.

We test the word representations on the 20 textual similarity datasets from the 2012-2015 SemEval semantic textual similarity (STS) shared tasks (Agirre et al., 2012, 2013, 2014; Agirre et al., 2015), and the 2012 SemEval Semantic Related task (SICK) (Marelli et al., 2014).

Representing sentences by the average of their constituent word representations is surprisingly effective in encoding the semantic information of sentences (Wieting et al., 2015; Adi et al., 2016) and close to the state-of-the-art in these datasets. We follow this rubric and represent a sentence s based on its averaged word representation, i.e.,

$v(s) = 1/|s| \sum_{w \in s} v(w)$, and then compute the similarity between two sentences via the cosine similarity between the two representations.

The detailed performance of the original and processed representations is itemized in Table 5 – we see a consistent and significant improvement in performance because of postprocessing (on average 4% improvement).

	WORD2VEC		GLOVE	
	orig.	proc.	orig.	proc.
2012	57.22	57.67	48.27	54.06
2013	56.81	57.98	44.83	57.71
2014	62.89	63.30	51.11	59.23
2015	62.74	63.35	47.23	57.29
SICK	70.10	70.20	65.14	67.85
all	60.88	61.45	49.19	56.76

Table 5: Before-After results (x100) on the semantic textual similarity tasks.

4 Neural Networks Learn to Postprocess

Supervised downstream NLP applications have greatly improved their performances in recent years by combining the discriminative learning powers of neural networks in conjunction with the word representations. All neural network architectures, ranging from feedforward to recurrent (either vanilla or GRU or LSTM), implement at least linear processing of hidden/input state vectors at each of their nodes; thus the postprocessing operation suggested in this paper can in principle be automatically “learnt” by the neural network, if such internal learning is in-line with the end-to-end training examples. Following this line of thought, the natural conjecture would be that both the original and postprocessed word representations should yield similar end-to-end results.

4.1 Postprocessing does not Affect Neural Network Performance

We evaluate the performance of a variety of neural network architectures on a standard and important NLP application: *text classification*, with sentiment analysis being a particularly important and popular example. The task is defined as follows: given a sentence, the algorithm needs to decide which category it falls into. The categories can be either binary (e.g., positive/negative) or can be more fine-grained (e.g. very positive, positive, neutral, negative, and very negative).

We evaluate the word representations (with and without postprocessing) using four different neural network architectures (CNN, vanilla-RNN, GRU-RNN and LSTM-RNN) on five benchmarks: (a) the movie review (MR) dataset (Pang and Lee, 2005); (b) the subjectivity (SUBJ) dataset (Pang and Lee, 2004); (c) the TREC question dataset (Li and Roth, 2002); (d) the IMDb dataset (Maas et al., 2011); (e) the stanford sentiment treebank (SST) dataset (Socher et al., 2013a). A detailed description of these standard datasets, their training/test parameters and the cross validation methods adopted is in Appendix B. The performance of the four neural network architectures with the now-standard CNN-based text classification algorithm (Kim, 2014) (implemented using tensorflow⁴) is itemized in Table 6 – the key observation is that the two performances, with and without postprocessing, are comparable (although postprocessing is mildly better in a majority of the instances). We hypothesize that this happens because the weight and bias in the neural network automatically encode the postprocessing; this is verified next.

4.2 An Appended-Layer Neural Network

We construct a modified neural network by explicitly adding a “postprocessing unit” as the first layer of the RNN architecture (as in Figure 2). Such an appended layer can be used to test the first step (i.e., remove the mean vector) of the postprocessing algorithm; setting up a corresponding test of the second step (nulling the top components) is unclear (see Appendix C for a discussion).

In the modified neural network, the input word vectors are now $v(w) - b$ instead of $v(w)$. Here b is a bias vector trained *jointly* with the rest of the neural network parameters. Note that this is only a relabeling of the parameters from the perspective of the RNN architecture: the nonlinear activation function of the node is now operated on $A(v(w) - b) + b' = Av(w) + (b' - Ab)$ instead of the previous $Av(w) + b'$. Let $b_{\text{proc.}}$ and $b_{\text{orig.}}$ be the inferred biases when using the processed and original word representations, respectively.

We itemize the cosine similarity between $b_{\text{proc.}} + \mu$ and $b_{\text{orig.}}$ in Table 7 for the 5 different datasets and 3 different neural network architectures. In each case, the cosine similarity is remark-




Figure 2: Time-expanded RNN architecture with an appended layer involving linear bias.

	vanilla		GRU		LSTM	
	W.	G.	W.	G.	W.	G.
MR	64.8	57.9	42.4	82.8	35.1	62.2
SUBJ	45.2	49.6	47.7	64.9	32.9	48.2
TREC	82.8	92.4	86.4	94.0	87.1	84.5
SST	94.8	88.9	75.3	93.1	73.1	73.2
IMDb	18.9	94.9	71.3	91.8	93.6	43.5

Table 7: The cosine similarity (x100) between $b_{\text{proc.}} + \mu$ and $b_{\text{orig.}}$, where W. and G. stand for WORD2VEC and GLOVE respectively.

ably large (on average 0.69, in 300 dimensions) – in other words, trained neural networks implicitly postprocess the word vectors nearly exactly as we proposed.

5 Discussion

The postprocessing operation is extremely simple and yet, quite effective across a very wide range of both intrinsic and extrinsic evaluations of the word representations. A key question, unaddressed thus far, is whether the dominant singular directions signify anything, and if so what those aspects could be. Our preliminary results are discussed next; we begin by noting that the mean vector is strongly correlated with the first principle component direction (cosine similarity of more than 0.8, in 300 dimensions).

5.1 Significance of Nulled Vectors

Consider the representation of the words as viewed in terms of the top D PCA coefficients $\alpha_\ell(w)$, for $1 \leq \ell \leq D$. We find that these few coefficients encode the *frequency* of the word to a significant degree; Figure 3 illustrates the relation between the $(\alpha_1(w), \alpha_2(w))$ and the unigram probability $p(w)$, where the correlation is geometrically visible.

⁴<https://github.com/dennybritz/cnn-text-classification-tf>

	CNN				vanilla-RNN				GRU-RNN				LSTM-RNN			
	WORD2VEC		GLOVE		WORD2VEC		GLOVE		WORD2VEC		GLOVE		WORD2VEC		GLOVE	
	orig.	proc.	orig.	proc.	orig.	proc.	orig.	proc.	orig.	proc.	orig.	proc.	orig.	proc.	orig.	proc.
MR	70.80	71.27	71.01	71.11	77.30	77.37	77.69	77.87	78.72	78.65	78.38	78.78	79.07	79.20	79.02	79.81
SUBJ	87.14	87.33	86.98	87.25	90.38	90.70	91.34	91.05	91.35	91.48	91.79	91.92	91.81	92.16	92.18	91.86
TREC	87.80	89.00	87.60	89.00	89.60	89.60	89.60	88.20	90.80	92.20	91.00	91.40	90.60	91.60	91.60	92.40
SST	38.46	38.33	38.82	37.83	41.54	41.90	42.58	43.57	44.57	44.16	43.76	43.94	44.62	45.43	44.16	44.89
IMDb	86.68	87.12	87.27	87.10	78.84	82.17	81.91	82.72	83.83	83.92	83.89	84.08	84.72	85.25	84.96	85.47

Table 6: Before-After results (x100) on the text classification task using CNN (Kim, 2014) and vanilla RNN, GRU-RNN and LSTM-RNN.




Figure 3: The top two PCA directions (i.e., $\alpha_1(w)$ and $\alpha_2(w)$) encode frequency.

5.2 Angular Asymmetry of Representations

A modern understanding of word representations involves either PMI-based (including word2vec (Mikolov et al., 2010; Levy and Goldberg, 2014) and GloVe (Pennington et al., 2014)) or CCA-based spectral factorization approaches. While CCA-based spectral factorization methods have long been understood from a probabilistic (i.e., generative model) view point (Browne, 1979; Hotelling, 1936) and recently in the NLP context (Stratos et al., 2015), a corresponding effort for the PMI-based methods has only recently been conducted in an inspired work (Arora et al., 2016).

Arora et al. (2016) propose a generative model (named RAND-WALK) of sentences, where every word is parameterized by a d -dimensional vector. With a key postulate that the word vectors are angularly uniform (“isotropic”), the family of PMI-based word representations can be explained under the RAND-WALK model in terms of the maximum likelihood rule.

We have observed that word vectors learnt through PMI-based approaches are not of zero-mean and are not isotropic (c.f. Section 2) contradicts with this postulate. The isotropy conditions are relaxed in Section 2.2 of (Arora et al., 2016), but the match with the spectral properties observed in Figure 1 is not immediate.

In this section, we resolve this by explicitly relaxing the constraints on the word vectors to directly fit the observed spectral properties. The relaxed conditions are: the word vectors should be isotropic around a point (whose distance to the origin is a small fraction of the average norm of word vectors) lying on a low dimensional subspace. Our main result is to show that even with this enlarged parameter-space, the maximum likelihood rule continues to be close to the PMI-based spectral factorization methods. A formal proof that integrates with the technical results of (Arora et al., 2016), and a discussion where we interpret the postprocessing operation as a “rounding” or “projection” towards isotropy (with better self-normalization properties (Andreas and Klein, 2015)) is in Appendix D.

6 Conclusion

We present a simple postprocessing operation that renders word representations even stronger. Due to their popularity, we have used the published representations of WORD2VEC and GLOVE in English in the main text of this paper; postprocessing continues to be successful for other representations and in multilingual settings – the detailed empirical results are tabulated in Appendix A.

References

- Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2016. Fine-grained analysis of sentence embeddings using auxiliary prediction tasks. *arXiv preprint arXiv:1608.04207*.
- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2014. Semeval-2014 task 10: Multilingual semantic textual similarity. In *Proceedings of the 8th international workshop on semantic evaluation (SemEval 2014)*. pages 81–91.
- Eneko Agirre, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, and Weiwei Guo. 2013. sem 2013 shared task: Semantic textual similarity, including a pilot on typed-similarity. In *In* SEM 2013: The Second Joint Conference on Lexical and Computational Semantics. Association for Computational Linguistics*. Citeseer.
- Eneko Agirre, Mona Diab, Daniel Cer, and Aitor Gonzalez-Agirre. 2012. Semeval-2012 task 6: A pilot on semantic textual similarity. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics-Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation*. Association for Computational Linguistics, pages 385–393.
- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Inigo Lopez-Gazpio, Montse Maritxalara, Rada Mihalcea, et al. 2015. Semeval-2015 task 2: Semantic textual similarity, english, spanish and pilot on interpretability. In *Proceedings of the 9th international workshop on semantic evaluation (SemEval 2015)*. pages 252–263.
- Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2013. Polyglot: Distributed word representations for multilingual nlp. *arXiv preprint arXiv:1307.1662*.
- Abdulrahman Almuhrareb. 2006. *Attributes in lexical acquisition*. Ph.D. thesis, University of Essex.
- Jacob Andreas and Dan Klein. 2015. When and why are log-linear models self-normalizing? In *HLT-NAACL*. pages 244–249.
- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. 2016. [A latent variable model approach to pmi-based word embeddings](https://arxiv.org/abs/1607.03903). *Transactions of the Association for Computational Linguistics* 4:385–399. <https://transacl.org/ojs/index.php/tacl/article/view/742>.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- M Baroni, S Evert, and A Lenci. 2008. Bridging the gap between semantic theory and computational simulations: Proceedings of the esslli workshop on distributional lexical semantics. *Hamburg, Germany: FOLLI*.
- Marco Baroni, Georgiana Dinu, and Germán Kruszewski. 2014. Don’t count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In *ACL (1)*. pages 238–247.
- Marco Baroni and Alessandro Lenci. 2010. Distributional memory: A general framework for corpus-based semantics. *Computational Linguistics* 36(4):673–721.
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. 2003. A neural probabilistic language model. *journal of machine learning research* 3(Feb):1137–1155.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems*. pages 2787–2795.
- Michael W Browne. 1979. The maximum-likelihood solution in inter-battery factor analysis. *British Journal of Mathematical and Statistical Psychology* 32(1):75–86.
- Elia Bruni, Nam-Khanh Tran, and Marco Baroni. 2014. Multimodal distributional semantics. *J. Artif. Intell. Res.(JAIR)* 49(1-47).
- José Camacho-Collados, Mohammad Taher Pilehvar, and Roberto Navigli. 2015. A framework for the construction of monolingual and cross-lingual word similarity datasets. In *ACL (2)*. pages 1–7.
- Junyoung Chung, Caglar Gülçehre, Kyunghyun Cho, and Yoshua Bengio. 2015. Gated feedback recurrent neural networks. In *ICML*. pages 2067–2075.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research* 12(Aug):2493–2537.
- Paramveer Dhillon, Jordan Rodu, Dean Foster, and Lyle Ungar. 2012. Two step cca: A new spectral method for estimating vector models of words. *arXiv preprint arXiv:1206.6403*.
- Lev Finkelstein, Evgeniy Gabilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppin. 2001. Placing search in context: The concept revisited. In *Proceedings of the 10th international conference on World Wide Web*. ACM, pages 406–414.

- Daniela Gerz, Ivan Vulić, Felix Hill, Roi Reichart, and Anna Korhonen. 2016. Simverb-3500: A large-scale evaluation set of verb similarity. *arXiv preprint arXiv:1608.00869*.
- Klaus Greff, Rupesh K Srivastava, Jan Koutník, Bas R Steunebrink, and Jürgen Schmidhuber. 2016. Lstm: A search space odyssey. *IEEE transactions on neural networks and learning systems*.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2016. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*.
- Harold Hotelling. 1936. Relations between two sets of variates. *Biometrika* 28(3/4):321–377.
- Eric H Huang, Richard Socher, Christopher D Manning, and Andrew Y Ng. 2012. Improving word representations via global context and multiple word prototypes. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*. Association for Computational Linguistics, pages 873–882.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*.
- Beatrice Laurent and Pascal Massart. 2000. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics* pages 1302–1338.
- Omer Levy and Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. In *Advances in neural information processing systems*. pages 2177–2185.
- Xin Li and Dan Roth. 2002. Learning question classifiers. In *Proceedings of the 19th international conference on Computational linguistics-Volume 1*. Association for Computational Linguistics, pages 1–7.
- Thang Luong, Richard Socher, and Christopher D Manning. 2013. Better word representations with recursive neural networks for morphology. In *CoNLL*. pages 104–113.
- Andrew L Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*. Association for Computational Linguistics, pages 142–150.
- Christopher D Manning, Prabhakar Raghavan, Hinrich Schütze, et al. 2008. *Introduction to information retrieval*, volume 1. Cambridge university press Cambridge.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. A sick cure for the evaluation of compositional distributional semantic models. In *LREC*. pages 216–223.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Tomas Mikolov, Martin Karafiát, Lukas Burget, Jan Cernocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *Inter-speech*. volume 2, page 3.
- Andriy Mnih and Geoffrey Hinton. 2007. Three new graphical models for statistical language modelling. In *Proceedings of the 24th international conference on Machine learning*. ACM, pages 641–648.
- Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the 42nd annual meeting on Association for Computational Linguistics*. Association for Computational Linguistics, page 271.
- Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd annual meeting on association for computational linguistics*. Association for Computational Linguistics, pages 115–124.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*. volume 14, pages 1532–43.
- Kira Radinsky, Eugene Agichtein, Evgeniy Gabrilovich, and Shaul Markovitch. 2011. A word at a time: computing word relatedness using temporal semantic analysis. In *Proceedings of the 20th international conference on World wide web*. ACM, pages 337–346.
- Herbert Rubenstein and John B Goodenough. 1965. Contextual correlates of synonymy. *Communications of the ACM* 8(10):627–633.
- Tobias Schnabel, Igor Labutov, David Mimno, and Thorsten Joachims. 2015. Evaluation methods for unsupervised word embeddings. In *Proc. of EMNLP*.
- Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Ng. 2013a. Reasoning with neural tensor networks for knowledge base completion. In *Advances in Neural Information Processing Systems*. pages 926–934.
- Richard Socher, Alex Perelygin, Jean Y Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013b. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the conference on empirical methods in natural language processing (EMNLP)*. Citeseer, volume 1631, page 1642.

- Karl Stratos, Michael Collins, and Daniel Hsu. 2015. Model-based word embeddings from decompositions of count matrices. In *Proceedings of ACL*, pages 1282–1291.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Joseph Turian, Lev Ratinov, and Yoshua Bengio. 2010. Word representations: a simple and general method for semi-supervised learning. In *Proceedings of the 48th annual meeting of the association for computational linguistics*. Association for Computational Linguistics, pages 384–394.
- John Wieting, Mohit Bansal, Kevin Gimpel, Karen Livescu, and Dan Roth. 2015. From paraphrase database to compositional paraphrase model and back. *arXiv preprint arXiv:1506.03487*.
- Torsten Zesch and Iryna Gurevych. 2006. Automatically creating datasets for measures of semantic relatedness. In *Proceedings of the Workshop on Linguistic Distances*. Association for Computational Linguistics, pages 16–24.

Appendix: All-but-the-Top: Simple and Effective postprocessing for Word Representations

A Experiments on Various Representations

In the main text, we have reported empirical results for two published word representations: WORD2VEC and GLOVE, each in 300 dimensions. In this section, we report the results of the same experiments in a variety of other settings to show the generalization capability of the postprocessing operation: representations trained via WORD2VEC and GLOVE algorithms in dimensions other than 300, other representations algorithms (specifically TSCCA and RAND-WALK) and in multiple languages.

A.1 Multilingual Generalization

In this section, we perform the word similarity task with the original and the processed TSCCA word representations in German and Spanish on three German similarity datasets (GUR65 – a German version of the RG65 dataset, GUR350, and ZG222 in terms of relatedness) (Zesch and Gurevych, 2006) and the Spanish version of RG65 dataset (Camacho-Collados et al., 2015). The choice of D is 2 for both German and Spanish.

The experiment setup and the similarity scoring algorithm are the same as those in Section 3.1. The detailed experiment results are provided in Table 8, from which we observe that the processed representations are consistently better than the original ones. This provides evidence to the generalization capabilities of the postprocessing operation to multiple languages (similarity datasets in Spanish and German were the only ones we could locate).

	language	TSCCA	
		orig.	proc.
RG65	Spanish	60.33	60.37
GUR65	German	61.75	64.39
GUR350	German	44.91	46.59
ZG222	German	30.37	32.92

Table 8: Before-After results (x100) on the word similarity task in multiple languages.

A.2 Generalization to Different Representation Algorithms

Given the popularity and widespread use of WORD2VEC (Mikolov et al., 2013) and GLOVE (Pennington et al., 2014), the main text has solely focused on their published publicly available 300-dimension representations. In this section, we show that the proposed postprocessing algorithm generalizes to other representation methods. Specifically, we demonstrate this on RAND-WALK (obtained via personal communication) and TSCCA (publicly available) on all the experiments of Section 3. The choice of D is 2 for both RAND-WALK and TSCCA.

In summary, the performance improvements on the similarity task, the concept categorization task, the analogy task, and the semantic textual similarity dataset are on average 2.23%, 2.39%, 0.11% and 0.61%, respectively. The detailed statistics are provided in Table 9, Table 10, Table 11 and Table 12, respectively. These results are a testament to the generalization capabilities of the postprocessing algorithm to other representation algorithms.

A.3 Role of Dimensions

The main text has focused on the dimension choice of $d = 300$, due to its popularity. In this section we explore the role of the dimension in terms of both choice of D and the performance of the postprocessing operation – we do this by using skip-gram model on the 2010 snapshot of Wikipedia corpus (Al-Rfou et al., 2013) to train word representations. We first observe that the two phenomena of Section 2 continue to hold:

	RAND-WALK		TSCCA	
	orig.	proc.	orig.	proc.
RG65	80.66	82.96	47.53	47.67
WS	65.89	74.37	54.21	54.35
RW	45.11	51.23	43.96	43.72
MEN	73.56	77.22	65.48	65.62
MTurk	64.35	66.11	59.65	60.03
SimLex	34.05	36.55	34.86	34.91
SimVerb	16.05	21.84	23.79	23.83

Table 9: Before-After results (x100) on the word similarity task on seven datasets.

	RAND-WALK		TSCCA	
	orig.	proc.	orig.	proc.
ap	59.83	62.36	60.00	63.42
esslli	72.73	72.73	68.18	70.45
battig	75.73	81.82	70.73	70.73

Table 10: Before-After results (x100) on the categorization task.

	RAND-WALK		TSCCA	
	orig.	proc.	orig.	proc.
syn.	60.39	60.48	37.72	37.80
sem.	83.55	83.82	14.54	14.55
all	70.50	70.67	27.30	27.35

Table 11: Before-After results (x100) on the word analogy task.

	RAND-WALK		TSCCA	
	orig.	proc.	orig.	proc.
2012	38.03	37.66	44.51	44.63
2013	37.47	36.85	43.21	42.74
2014	46.06	48.32	52.85	52.87
2015	47.82	51.76	56.22	56.14
SICK	51.58	51.76	56.15	56.11
all	43.48	44.67	50.01	50.23

Table 12: Before-After results (x100) on the semantic textual similarity tasks.

- From Table 13 we observe that the ratio between the norm of μ and the norm average of all $v(w)$ spans from 1/3 to 1/4;

dim	300	400	500	600	700	800	900	1000
avg. $\ v(w)\ _2$	4.51	5.17	5.91	6.22	6.49	6.73	6.95	7.15
$\ \mu\ _2$	1.74	1.76	1.77	1.78	1.79	1.80	1.81	1.83

Table 13: Statistics on word representation of dimensions 300, 400, ..., and 1000 using the skip-gram model.

- From Figure 4 we observe that the decay of the variance ratios σ_i is near exponential for small values of i and remains roughly constant over the later ones.

A rule of thumb choice of D is around $d/100$. We validate this claim empirically by performing the tasks in Section 3 on word representations of higher dimensions, ranging from 300 to 1000, where we




Figure 4: The decay of the normalized singular values of word representations.

set the parameter $D = d/100$. In summary, the performance improvement on the four itemized tasks of Section 3 are 2.27%, 3.37%, 0.01 and 1.92% respectively; the detailed results can be found in Table 14, Table 15, Table 16, and Table 17. Again, note that the improvement for analogy tasks is marginal. These experimental results justify the rule-of-thumb setting of $D = d/100$, although we emphasize that the improvements can be further accentuated by tuning the choice of D based on the specific setting.

Dim	300		400		500		600	
	orig.	proc.	orig.	proc.	orig.	proc.	orig.	proc.
RG65	73.57	74.72	75.64	79.87	77.72	81.97	77.59	80.7
WS	70.25	71.95	70.8	72.88	70.39	72.73	71.64	74.04
RW	46.25	49.11	45.97	47.63	46.6	48.59	45.7	47.81
MEN	75.66	77.59	76.07	77.89	75.9	78.15	75.88	78.15
Mturk	75.66	77.59	67.68	68.11	66.89	68.25	67.6	67.87
SimLex	34.02	36.19	35.17	37.1	35.73	37.65	35.76	38.04
SimVerb	22.22	24.98	22.91	25.32	23.03	25.82	23.35	25.97
Dim	700		800		900		1000	
	orig.	proc.	orig.	proc.	orig.	proc.	orig.	proc.
RG65	77.3	81.07	77.52	81.07	79.75	82.34	78.18	79.07
WS	70.31	73.02	71.52	74.65	71.19	73.06	71.5	74.78
RW	45.86	48.4	44.96	49	44.44	49.22	44.5	49.03
MEN	75.84	78.21	75.84	77.96	76.16	78.35	76.72	78.1
Mturk	67.47	67.79	67.67	68	67.98	68.87	68.34	69.44
SimLex	35.3	37.59	36.54	37.85	36.62	38.44	36.67	38.58
SimVerb	22.81	25.6	23.48	25.57	23.68	25.76	23.24	26.58

Table 14: Before-After results (x100) on word similarity task on seven datasets.

B Statistics of Text Classification Datasets

We evaluate the word representations (with and without postprocessing) using four different neural network architectures (CNN, vanilla-RNN, GRU-RNN and LSTM-RNN) on five benchmarks:

- the movie review (MR) dataset (Pang and Lee, 2005) where each review is composed by only one sentence;
- the subjectivity (SUBJ) dataset (Pang and Lee, 2004) where the algorithm needs to decide whether a sentence is subjective or objective;

Dim	300		400		500		600	
	orig.	proc.	orig.	proc.	orig.	proc.	orig.	proc.
ap	46.1	48.61	42.57	45.34	46.85	50.88	40.3	45.84
esslli	68.18	72.73	64.2	82.72	64.2	65.43	65.91	72.73
battig	71.6	77.78	68.18	75	68.18	70.45	46.91	66.67
Dim	700		800		900		1000	
	orig.	proc.	orig.	proc.	orig.	proc.	orig.	proc.
ap	38.04	41.31	34.76	39.8	34.76	27.46	27.96	28.21
esslli	54.55	54.55	68.18	56.82	72.73	72.73	52.27	52.27
battig	62.96	66.67	67.9	69.14	49.38	59.26	51.85	46.91

Table 15: Before-After results (x100) on the categorization task.

Dim	300		400		500		600	
	orig.	proc.	orig.	proc.	orig.	proc.	orig.	proc.
syn.	60.48	60.52	61.61	61.45	60.93	60.84	61.66	61.57
sem.	74.51	74.54	77.11	77.36	76.39	76.89	77.28	77.61
all.	66.86	66.87	68.66	68.69	67.88	68.11	68.77	68.81
Dim	700		800		900		1000	
	orig.	proc.	orig.	proc.	orig.	proc.	orig.	proc.
syn.	60.94	61.02	68.38	68.34	60.47	60.30	67.56	67.30
sem.	77.24	77.26	77.24	77.35	76.76	76.90	76.71	76.51
all.	68.36	68.41	68.38	68.50	67.91	67.67	67.56	67.30

Table 16: Before-After results (x100) on the word analogy task.

Dim	300		400		500		600	
	orig.	proc.	orig.	proc.	orig.	proc.	orig.	proc.
2012	54.51	54.95	54.31	54.57	55.13	56.23	55.35	56.03
2013	56.58	57.89	56.35	57.35	57.55	59.38	57.43	59.00
2014	59.6	61.92	59.57	61.62	61.19	64.38	61.10	63.86
2015	59.65	61.48	59.69	61.19	61.63	64.77	61.42	64.04
SICK	68.89	70.79	60.6	70.27	68.63	71.00	68.58	70.57
all	58.32	59.91	58.25	59.55	59.61	62.02	59.57	61.55
Dim	700		800		900		1000	
	orig.	proc.	orig.	proc.	orig.	proc.	orig.	proc.
2012	55.52	56.49	54.47	54.85	54.69	55.18	54.34	54.78
2013	57.61	59.31	56.75	57.62	56.98	58.26	56.78	57.73
2014	61.57	64.77	60.51	62.83	60.89	63.34	60.78	63.03
2015	62.05	65.45	60.74	62.84	61.09	63.48	60.92	63.03
SICK	68.38	70.63	67.94	69.59	67.86	69.5	67.58	69.16
all	59.96	62.34	58.87	60.39	59.16	60.88	58.94	60.48

Table 17: Before-After results (x100) on the semantic textual similarity tasks.

- the TREC question dataset (Li and Roth, 2002) where all the questions in this dataset has to be partitioned into six categories;
- the IMDb dataset (Maas et al., 2011) – each review consists of several sentences;
- the Stanford sentiment treebank (SST) dataset (Socher et al., 2013a), where we only use the full sentences as the training data.

In TREC, SST and IMDB, the datasets have already been split into train/test sets. Otherwise we use 10-fold cross validation in the remaining datasets (i.e., MR and SUBJ). Detailed statistics of various features of each of the datasets are provided in Table 18.

	c	l	Train	Test
MR	2	20	10,662	10-fold cross validation
SUBJ	2	23	10,000	10-fold cross validation
TREC	6	10	5,952	500
SST	5	18	11,855	2,210
IMDb	2	100	25,000	25,000

Table 18: Statistics for the five datasets after tokenization: c represents the number of classes; l represents the average sentence length; Train represents the size of the training set; and Test represent the size of the test set.

C Neural Networks and Postprocessed Word Representations

Every neural network family possesses the ability to conduct linear processing inside their nodes; this includes feedforward and recurrent and convolutional neural network models. Thus, in principle, the postprocessing operation can be “learnt and implemented” within the parameters of the neural network. On the other hand, due to the large number of parameters within the neural network, it is unclear how to verify such a process, even if it were learnt (only one of the layers might be implementing the postprocessing operation or via a combination of multiple effects).

To address this issue, we have adopted a *comparative* approach in the main text of the paper. The comparative approach involves adding an *extra layer* interposed in between the inputs (which are word vectors) and the rest of the neural network. This extra layer involves only linear processing. Next we compare the results of the final parameters of the extra layer (trained jointly with the rest of the neural network parameters, using the end-to-end training examples) with and without preprocessing of the word vectors. Such a comparative approach allows us to separate the effect of the postprocessing operation on the word vectors from the complicated “semantics” of the neural network parameters.

This agenda is successfully implemented in the context of verifying the removal of the mean vector (and is described in the main text). A similar effort for the second step (nulling the dominant PCA directions) is not possible with this approach: this is because the nulling operation cannot be verified by a comparison of the (matrix) weights of the appended layer. The nulling would have to be verified solely by checking if the appended layer would null the dominant PCA directions *directly* on the original word vectors. – this test is perhaps too strong a requirement, as discussed above. An experimental setup to verify the second step of the postprocessing operation is left as a future research direction.

D Angular Asymmetry of Representations

In the main text, we have observed that the word vectors learnt through PMI-based approaches are not of zero-mean and are not isotropic (c.f. Section 2) contradicts with the Bayesian prior postulate in (Arora et al., 2016). Such a contradiction suggests that either there is a mismatch between the RAND-WALK model and the text data, or the isotropy postulate is too restrictive (or both). The conditions on isotropy of the word vectors are somewhat relaxed in Sections 2.2 and 4 of (Arora et al., 2016), but the match with the spectral properties observed in Figure 1 is not immediate.

In this section, we resolve this mismatch by directly relaxing the constraints on the word vectors to match the observed spectral properties. Specifically, the word vectors should be isotropic when projected over a subspace (which itself is another parameter) of large enough dimension. Our main result is to show that even with this enlarged parameter-space, the maximum likelihood rule continues to be close to the PMI-based spectral factorization methods; we do this by building upon the technical foundation laid in (Arora et al., 2016).

Formally, the model, the original constraints of (Arora et al., 2016) and the enlarged constraints on the word vectors are listed below:

- **A generative model of sentences:** the word at time t , denoted by w_t , is generated via a log-linear model with a latent discourse variable c_t (Arora et al., 2016), i.e.,

$$p(w_t|c_t) = \frac{1}{Z(c_t)} \exp(c_t^T v(w_t)), \quad (1)$$

where $v(w) \in \mathbb{R}^d$ is the vector representation for a word w in the vocabulary V , c_t is the latent variable which forms a “slowly moving” random walk, and the partition function: $Z(c) = \sum_{w \in V} \exp(c^T v(w))$.

- **Constraints on the word vectors:** Arora et al. (2016) suppose that there is a Bayesian prior on the word vectors:

The ensemble of word vectors consists of i.i.d. draws generated by $v = s \cdot \hat{v}$, where \hat{v} is from the spherical Gaussian distribution, and s is a scalar random variable.

A deterministic version of this Bayesian prior is discussed in Section 2.2 of (Arora et al., 2016), but part of these (relaxed) conditions on the word vectors are specifically meant for Theorem 4.1 and not the main theorem (Theorem 2.2). The geometry of the word representations is only evaluated via the ratio of the quadratic *mean* of the singular values to the smallest one being small enough. This meets the relaxed conditions, but not sufficient to validate the proof approach of the main result (Theorem 2.2); what would be needed is that the ratio of the *largest* singular value to the smallest one be small.

- **Revised conditions:** Below we revise the Bayesian prior postulate (and in a deterministic fashion) formally as follows: there is a mean vector μ , D orthonormal vectors u_1, \dots, u_D (that are orthogonal and of unit norm), such that every word vector $v(w)$ can be represented by,

$$v(w) = \mu + \sum_{i=1}^D \alpha_i(w) u_i + \tilde{v}(w), \quad (2)$$

where μ is bounded, α_i is bounded by A , D is bounded by $DA^2 = o(d)$, $\tilde{v}(w)$ are statistically isotropic. By statistical isotropy, we mean: for high-dimensional rectangles R , $\frac{1}{|V|} \sum_{w \in V} \mathbf{1}(\tilde{v}(w) \in R) \rightarrow \int_R f(\tilde{v}) d\tilde{v}$, as $|V| \rightarrow \infty$, where f is an angle-independent pdf, i.e., $f(\tilde{v})$ is a function of $\|\tilde{v}\|$.

The revised postulate differs from the original one in two ways: (a) it imposes a formal deterministic constraint on the word vectors; (b) the revised postulate allows the word vectors to be angularly asymmetric: as long as the energy in the direction of u_1, \dots, u_D is bounded, there is no constraint on the coefficients. Indeed, note that there is no constraint on $\tilde{v}(w)$ to be orthogonal to u_1, \dots, u_D .

Empirical Validation We can verify that the enlarged conditions are met by the existing word representations. Specifically, the natural choice for μ is the mean of the word representations and $u_1 \dots u_D$ are the singular vectors associated with the top D singular values of the matrix of word vectors. We pick $D = 20$ for WORD2VEC and $D = 10$ for GLOVE, and the corresponding value of DA^2 for WORD2VEC and GLOVE vectors are both roughly 40, respectively; both values are small compared to $d = 300$.

This leaves us to check the statistical isotropy of the “remaining” vectors $\tilde{v}(w)$ for words w in the vocabulary. We do this by plotting the remaining spectrum (i.e. the $(D + 1)$ -th, ..., 300th singular values) for the published WORD2VEC and GLOVE vectors in Figure 5. As a comparison, the empirical spectrum of a random Gaussian matrix is also plotted in Figure 5. We see that both spectra are flat (since the vocabulary size is much larger than the dimension $d = 300$). Thus the postprocessing operation can also be viewed as a way of making the vectors “more isotropic”. An alternative way to verify this is to check if the “self-normalization” property (i.e., $Z(c)$ is a constant, independent of c) holds more strongly. Such a validation is seen diagrammatically in Figure 6 where we randomly sampled 1,000 c ’s as (Arora et al., 2016)




Figure 5: Spectrum of the published WORD2VEC and GLOVE and random Gaussian matrices, ignoring the top D components; $D = 10$ for GLOVE and $D = 20$ for WORD2VEC.




Figure 6: The histogram of $Z(c)$ for 1,000 randomly sampled vectors c of unit norm, where the x -axis is normalized by the mean of all values, where we choose $D = 10$ for GLOVE and $D = 20$ for WORD2VEC.

Mathematical Contribution The revised postulate, we show that the main theorem in (Arora et al., 2016) (c.f. Theorem 2.2) still holds. Formally:

Theorem D.1 *Suppose the word vectors satisfy the constraints. Then*

$$\text{PMI}(w_1, w_2) \stackrel{\text{def}}{=} \log \frac{p(w_1, w_2)}{p(w_1)p(w_2)} \rightarrow \frac{v(w_1)^T v(w_2)}{d}, \quad \text{as } |V| \rightarrow \infty, \quad (3)$$

where $p(w)$ is the unigram distribution induced from the model (1), and $p(w_1, w_2)$ is the probability that two words w_1 and w_2 occur with each other within distance q .

Theorem D.1 suggests that the RAND-WALK generative model and its properties proposed by (Arora et al., 2016) can be generalized to a broader setting (with a relaxed restriction on the geometry of word representations) – relevantly, this relaxation on the geometry of word representations is empirically satisfied by the vectors learnt as part of the maximum likelihood rule.

Given the similarity between the setup in Theorem 2.2 in (Arora et al., 2016) and Theorem D.1, many parts of the original proof can be reused except one key aspect – the concentration of $Z(c)$. We summarize this part in the following lemma:

Lemma D.2 *Let c be a random variable uniformly distributed over the unit sphere, we prove that with high probability, $Z(c)/|V|$ converges to a constant Z :*

$$p((1 - \epsilon_z)Z \leq Z(c) \leq (1 + \epsilon_z)Z) \geq 1 - \delta,$$

where $\epsilon_z = \Omega((D + 1)/|V|)$ and $\delta = \Omega((DA^2 + \|\mu\|^2)/d)$.

Our proof differs from the one in (Arora et al., 2016) in two ways: (a) we treat $v(w)$ as deterministic parameters instead of random variables and prove the Lemma by showing a certain concentration of measure; (b) the asymmetric parts μ and u_1, \dots, u_D , (which did not exist in the original proof), need to be carefully addressed to complete the proof.

D.1 Proof for Lemma D.2

Given the constraints on the word vectors (2), the partition function $Z(c)$ can be rewritten as,

$$\begin{aligned} Z(c) &= \sum_{v \in V} \exp(c^T v(w)) \\ &= \sum_{v \in V} \exp \left(c^T \left(\mu + \sum_{i=1}^D \alpha_i(w) u_i + \tilde{v}(w) \right) \right) \\ &= \sum_{v \in V} \exp(c^T \mu) \left[\prod_{i=1}^D \exp(\alpha_i(w) c^T u_i) \right] \exp(c^T \tilde{v}(w)). \end{aligned}$$

The equation above suggests that we can divide the proof into five parts.

Step 1: for every unit vector c , one has,

$$\frac{1}{|V|} \sum_{w \in V} \exp(c^T \tilde{v}(w)) \rightarrow \mathbb{E}_f(\exp(c^T \tilde{v})), \text{ as } |V| \rightarrow \infty. \quad (4)$$

Proof Let M, N be a positive integer, and let $A_M \subset \mathbb{R}^d$ such that,

$$A_{M,N} = \left\{ \tilde{v} \in \mathbb{R}^d : \frac{M-1}{N} < \exp(c^T \tilde{v}) \leq \frac{M}{N} \right\}.$$

Since $A_{M,N}$ can be represented by a union of countable disjoint rectangles, we know that for every $M, N \in \mathbb{N}_+$,

$$\frac{1}{|V|} \sum_{w \in V} \mathbf{1}(\tilde{v}(w) \in A_{M,N}) = \int_{A_{M,N}} f(\tilde{v}) d\tilde{v}.$$

Further, since $A_{M,N}$ are disjoint for different M 's and $\mathbb{R}^d = \cup_{M=1}^{\infty} A_{M,N}$, one has,

$$\begin{aligned} \frac{1}{|V|} \sum_{w \in V} \exp(c^T \tilde{v}(w)) &= \sum_{M=1}^{\infty} \frac{1}{|V|} \sum_{w \in V} \mathbf{1}(\tilde{v}(w) \in A_{M,N}) \exp(c^T \tilde{v}(w)) \\ &\leq \sum_{M=1}^{\infty} \frac{1}{|V|} \sum_{w \in V} \mathbf{1}(\tilde{v}(w) \in A_{M,N}) \frac{M}{N} \\ &\rightarrow \sum_{M=1}^{\infty} \frac{M}{N} \int_{A_{M,N}} f(\tilde{v}) d\tilde{v}. \end{aligned}$$

The above statement holds for every N . Let $N \rightarrow \infty$, by definition of integration, one has,

$$\lim_{N \rightarrow \infty} \sum_{M=1}^{\infty} \frac{M}{N} \int_{A_{M,N}} f(\tilde{v}) d\tilde{v} = \mathbb{E}_f(\exp(c^T \tilde{v})),$$

which yields,

$$\frac{1}{|V|} \sum_{w \in V} \exp(c^T \tilde{v}(w)) \leq \mathbb{E}_f(\exp(c^T \tilde{v})), \text{ as } |V| \rightarrow \infty. \quad (5)$$

Similarly, one has,

$$\begin{aligned} \frac{1}{|V|} \sum_{w \in V} \exp(c^T \tilde{v}(w)) &\geq \lim_{N \rightarrow \infty} \sum_{M=1}^{\infty} \frac{M-1}{N} \int_{A_{M,N}} f(\tilde{v}) d\tilde{v} \\ &= \mathbb{E}_f(\exp(c^T \tilde{v})), \text{ as } |V| \rightarrow \infty. \end{aligned} \quad (6)$$

Putting (5) and (6) proves (4).

Step 2: the expected value, $\mathbb{E}_f(\exp(c^T \tilde{v}))$ is a constant independent of c :

$$\mathbb{E}_f(\exp(c^T \tilde{v})) = Z_0. \quad (7)$$

Proof Let $Q \in \mathbb{R}^{d \times d}$ be a orthonormal matrix such that $Q^T c_0 = c$ where $c_0 = (1, 0, \dots, 0)^T$ and $\det(Q) = 1$, then we have $f(\tilde{v}) = f(Q\tilde{v})$, and,

$$\begin{aligned} \mathbb{E}_f(\exp(c_0^T \tilde{v})) &= \int_{\tilde{v}} f(\tilde{v}) \exp(c_0^T \tilde{v}) d\tilde{v} \\ &= \int_{\tilde{v}} f(Q\tilde{v}) \exp(c^T Q\tilde{v}) \det(Q) d\tilde{v} \\ &= \int_{\tilde{v}'} f(\tilde{v}') \exp(c^T \tilde{v}') d\tilde{v}' = \mathbb{E}_f(\exp(c^T \tilde{v})), \end{aligned}$$

which proves (7).

Step 3: for any vector μ , one has the following concentration property,

$$p(|\exp(c^T \mu) - 1| > k) \leq 2 \left[\exp\left(-\frac{1}{4}\right) + \frac{\|\mu\|^2}{d-1} \frac{1}{\log^2(1-k)} \right] \quad (8)$$

Proof Let c_1, \dots, c_d be i.i.d. $\mathcal{N}(0, 1)$, and let $C = \sum_{i=1}^d c_i^2$, then $c = (c_1, \dots, c_d)/\sqrt{C}$ is uniform over unit sphere. Since c is uniform, then without loss of generality we can consider $\mu = (\|\mu\|, 0, \dots, 0)$. Thus it suffices to bound $\exp(\|\mu\|c_1/\sqrt{C})$. We divide the proof into the following steps:

- C follows chi-square distribution with the degree of freedom of d , thus C can be bounded by (Laurent and Massart, 2000),

$$p(C \geq d + 2\sqrt{dx} + 2x) \leq \exp(-x), \forall x > 0. \quad (9)$$

$$p(C \leq d - 2\sqrt{dx}) \leq \exp(-x), \forall x > 0. \quad (10)$$

- Therefore for any $x > 0$, one has,

$$p(|C - d| \geq 2\sqrt{dx}) \leq \exp(-x)$$

Let $x = 1/4d$, one has,

$$\begin{aligned} p(C > d + 1) &\leq \exp\left(-\frac{1}{4d}\right), \\ p(C < d - 1) &\leq \exp\left(-\frac{1}{4d}\right). \end{aligned}$$

- Since c_1 is a Gaussian random variable with variance 1, by Chebyshev's inequality, one has,

$$p(y c_i \geq k) \leq \frac{y^2}{k^2}, \quad p(y c_i \leq -k) \leq \frac{y^2}{k^2}, \forall k > 0$$

and therefore thus,

$$p(\exp(y c_i) - 1 > k) \leq \frac{y^2}{\log^2(1+k)},$$

$$p(\exp(y c_i) - 1 < -k) \leq \frac{y^2}{\log^2(1-k)^2}, \forall k > 0.$$

- Therefore we can bound $\exp\left(\frac{\|\mu\|_{c_1}}{\sqrt{C}}\right)$ by,

$$\begin{aligned} p\left(\exp\left(\frac{\|\mu\|_{c_1}}{\sqrt{C}}\right) - 1 > k\right) &\leq p(C > d+1) \\ &\quad + p\left(\exp\left(\frac{\|\mu\|_{c_1}}{\sqrt{C}}\right) - 1 > k \mid C < d+1\right) p(C < d+1) \\ &\leq \exp\left(-\frac{1}{4d}\right) + p\left(\exp\left(\frac{\|\mu\|_{c_1}}{\sqrt{d+1}}\right) - 1 > k\right) \\ &= \exp\left(-\frac{1}{4d}\right) + \frac{\|\mu\|^2}{d+1} \frac{1}{\log^2(1-k)^2}. \\ p\left(\exp\left(\frac{\|\mu\|_{c_1}}{\sqrt{C}}\right) - 1 < -k\right) &\leq \exp\left(-\frac{1}{4d}\right) + \frac{\|\mu\|^2}{d-1} \frac{1}{\log^2(1+k)^2}. \end{aligned}$$

Combining the two inequalities above, one has (8) proved.

Step 4: We are now ready to prove convergence of $Z(c)$. With (8), let $\mathcal{C} \subset \mathbb{R}^d$ such that,

$$\mathcal{C} = \{c : |\exp(c^\top \mu) - 1| < k, |\exp(Ac^\top u_i) - 1| < k, |\exp(-Ac^\top u_i) - 1| < k \forall i = 1, \dots, D\}$$

Then we can bound the probability on \mathcal{C} by,

$$\begin{aligned} p(\mathcal{C}) &\geq p(|\exp(c^\top \mu) - 1| < k) + \sum_{i=1}^D p(|\exp(Ac^\top u_i) - 1| < k) - 2D \\ &\geq 1 - (2D+1) \exp\left(-\frac{1}{4d}\right) - \frac{2DA^2}{d-1} \frac{1}{\log^2(1-k)} - \frac{\|\mu\|^2}{d-1} \frac{1}{\log^2(1-k)}. \end{aligned}$$

Next, we need to show that for every w , the corresponding $\mathcal{C}(w)$, i.e.,

$$\mathcal{C}(w) = \{c : |\exp(c^\top \mu) - 1| < k, |\exp(\alpha_i(w) c^\top u_i) - 1| < k, \forall i = 1, \dots, D\}$$

We observe that $\alpha_i(w)$ is bounded by A , therefore for any c that,

$$\min(\exp(-Ac^\top u_i), \exp(Ac^\top u_i)) \leq \exp(\alpha_i c^\top u_i) \leq \max(\exp(-Ac^\top u_i), \exp(Ac^\top u_i)),$$

and thus,

$$\min(\exp(-Ac^\top u_i), \exp(Ac^\top u_i)) - 1 \leq \exp(\alpha_i c^\top u_i) - 1 \leq \max(\exp(-Ac^\top u_i), \exp(Ac^\top u_i)) - 1,$$

which yields,

$$|\exp(\alpha_i c^\top u_i) - 1| \leq \max(|\exp(-Ac^\top u_i) - 1|, |\exp(Ac^\top u_i) - 1|) < k.$$

Therefore we prove $\mathcal{C}(w) \supset \mathcal{C}$. Assembling everything together, one has,

$$\begin{aligned} p\left(\left|\exp(c^\top \mu) \prod_{i=1}^D \exp(\alpha_i(w) c^\top u_i) - 1\right| > (D+1)k, \forall i = 1, \dots, D, \forall w \in V\right) \\ \leq p(\bar{\mathcal{C}}) \\ \leq (2D+1) \exp\left(-\frac{1}{4d}\right) + \frac{2DA^2}{d-1} \frac{1}{\log^2(1-k)} + \frac{\|\mu\|^2}{d-1} \frac{1}{\log^2(1-k)} \end{aligned}$$

For every $c \in \mathcal{C}$, one has,

$$\frac{1}{|V|} |Z(c) - Z_0| \leq \frac{(D+1)k}{|V|} Z_0.$$

Let $Z = |V|Z_0$, one can conclude that,

$$p((1 - \epsilon_z)Z \leq Z(c) \leq (1 + \epsilon_z)Z) \geq 1 - \delta,$$

where $\epsilon_z = \Omega((D+1)/|V|)$ and $\delta = \Omega(DA^2/d)$.

D.2 Proof for Theorem D.1

Having Lemma D.2 ready, we can follow the same proof as in (Arora et al., 2016) that both $p(w)$ and $p(w, w')$ are correlated with $\|v(w)\|$, formally

$$\log p(w) \rightarrow \frac{\|v(w)\|^2}{2d} - \log Z, \text{ as } |V| \rightarrow \infty, \quad (11)$$

$$\log p(w, w') \rightarrow \frac{\|v(w) + v(w')\|^2}{2d} - \log Z, \text{ as } |V| \rightarrow \infty. \quad (12)$$

Therefore, the inference presented in (Arora et al., 2016) (i.e., (3)) is obvious by assembling (11) and (12) together:

$$\text{PMI}(w, w') \rightarrow \frac{v(w)^T v(w')}{d}, \text{ as } |V| \rightarrow \infty.$$